

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06F 15/16	A1	(11) International Publication Number: WO 00/67142 (43) International Publication Date: 9 November 2000 (09.11.00)
<p>(21) International Application Number: PCT/US00/12201</p> <p>(22) International Filing Date: 4 May 2000 (04.05.00)</p> <p>(30) Priority Data: 09/305,427 5 May 1999 (05.05.99) US</p> <p>(71) Applicant: THE PORTOLA GROUP, INC. [US/US]; 3000 Sand Hill Road, Suite 2-145, Menlo Park, CA 94025 (US).</p> <p>(72) Inventors: MAKARIOS, Selene, K.; 900 High School Way, #2202, Mountain View, CA 94041 (US). FITZWILSON, Robert, C.; 130 Tuscaloosa Avenue, Atherton, CA 94025 (US). DOWNS, Heather, L.; 900 High School Way, #2202, Mountain View, CA 94041 (US).</p> <p>(74) Agents: BENGTTSSON, W., Patrick et al.; Pillsbury Madison & Sutro LLP, 50 Fremont Street, San Francisco, CA 94105 (US).</p>		<p>(81) Designated States: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published With international search report.</p>
<p>(54) Title: METHOD FOR COORDINATING ACTIVITIES AND SHARING INFORMATION USING A DATA DEFINITION LANGUAGE</p> <p>(57) Abstract</p> <p>A widely-used data definition language such as the Extensible Markup Language is used to implement a tuple space-based coordination mechanism. Entries and template entries can represent any type of networked or network-proxied resource, object or service. Using this framework, diverse entry spaces can be aggregated and operated upon as though they were a single large entry space. The flexibility and power of XML constructs can be leveraged to make such aggregation straightforward and efficient.</p> <p>BEST AVAILABLE COPY</p>		

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

METHOD FOR COORDINATING ACTIVITIES
AND SHARING INFORMATION
USING A DATA DEFINITION LANGUAGE

5 BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention is directed to a system and method for coordinating activities and sharing information between networked software entities. More specifically, the invention is directed towards implementing data exchange for such purposes using a data
10 definition language.

2. Background of the Related Art

The continuing development of parallel and distributed computing models brings with it issues of how to efficiently share information between networked or linked
15 processors. One possible solution to this problem is Linda, a coordination language (as opposed to a computation language such as C or FORTRAN) proposed by Nicholas Carriero and David Gelernter of Yale University. Linda is based on a logically associative object memory model called a tuple space. A tuple space is a virtual shared memory model which provides interprocess communication and synchronization logically independent of the
20 underlying computer system or network on which it resides. It uses a small number of simple operations on the tuple space to create and coordinate parallel processes which can simply exchange information.

The tuple space is so named because the fundamental data structure which populates it is called a tuple. A tuple is an ordered set of one or more fields with values, e.g.,
25 (sweater, wool, xxl-tall). Linda may use pattern-tuples, which are partially-specified

tuples denoted with wild cards. e.g., ?sweater, wool, ... There are four basic operations on tuples in tuple space:

-- for tuple generation, the operation `out(tuple)` is a non-blocking call (it does not stall program execution) that generates a data tuple (a tuple having static data) having specified values and puts it into tuple space. Control then returns to the invoking program. For example, the operation call

```
out(shirt, cotton, med)
```

puts a tuple (shirt, cotton, med) into a given tuple space.

-- also for tuple generation, the operation `eval(tuple)` generates a process tuple (a tuple under active evaluation) and returns. For arguments of the tuple which are a function call, conceptually processes are created to evaluate the functions. The results returned by the functions are substituted for the function calls in the tuple, and the tuple is placed into tuple space. For example, the operation call

```
eval('inventory', ii, inventory(ii))
```

might create a tuple which calls the function `inventory` to inventory the number of tuples having a certain field value of `i`. The result would be a data tuple in the tuple space which has the inventory result as its last field value.

-- for tuple extraction, the operation `in(pattern-tuple)` is a blocking call (it may stall program execution) which uses `pattern-tuple` to retrieve a tuple from the tuple space. The tuple is removed from the tuple-space and is no longer available for use by other processes. If no matching tuple is in the tuple space, the operation will stall until one becomes available. For example, the operation call

```
in(pants, ?material, large)
```

might look for all "large pants" tuples regardless of the material from which they are made. When and if it matches a tuple in tuple space, the value of the tuple's middle field will be assigned to material.

-- also for tuple extraction, the operation `rd(pattern-tuple)` is a blocking call that retrieves a copy of a matching tuple from the tuple space but leaves the original tuple in the tuple space. Thus, it may be thought of as a non-destructive version of `in`.

The Linda paradigm is both powerful and elegant, subsuming and organizing in a clean way some important issues in parallel and distributed computation and coordination. It has demonstrated its utility in a variety of applications involving coordinated software entities. For example, Linda forms the basis for a number of network device attachment and operation models. In particular, Linda features have been incorporated into several Java-based systems such as JavaSpaces, which forms a part of Sun Microsystems' Jini system; T Spaces from IBM Almaden Research laboratories; Java Paradise from Scientific Computing Associates, and Jada from the University of Bologna.

Unfortunately, these various Java implementations of Linda are not compatible with one another in the particulars of the entries which they store in their respective tuple spaces. This diminishes the potential for harmonious interoperation between systems speaking the various different dialects of Linda. Even if the different derivative implementations were mutually compatible, they all assume a language-based tuple space of Java objects and are not compatible with, e.g., non-Java based implementations.

SUMMARY OF THE INVENTION

With the above problems of the prior art in mind, it is an object of the present invention to provide a system and method for coordinating activities and sharing information among networked software entities.

5 It is a further object of the present invention to provide a unifying framework and standard for coordination and information sharing among networked software entities.

It is yet another object of the present invention to provide a system and method for combining networked tuple-spaces into larger distributed tuple-spaces without limit by exploiting the uniform representation which all tuple spaces share.

10 The above objects are achieved according to a first aspect of the invention by applying a widely-used data definition language such as the Extensible Markup Language (XML) to the domain of tuple space-based coordination mechanisms. With XML, for example, entries and template entries (similar to tuples and pattern tuples in Linda) are instances of XML Document Type Definitions (DTDs). These entries can represent any type
15 of networked or network-proxied resource, object or service. Using this framework, diverse entry spaces can be aggregated and operated upon as though they were a single large entry space. The flexibility and power of XML constructs can be leveraged to make such aggregation straightforward and efficient.

20

BRIEF DESCRIPTION OF THE DRAWINGS

These and other objects, features, and advantages of the present invention are better understood by reading the following detailed description of the preferred embodiment, taken in conjunction with the accompanying drawings, in which:

FIGURE 1 is a graph showing sub-linear scaling of element types with respect
25 to element number; and

FIGURE 2 is a tree diagram showing the structure of an element universe according to a preferred embodiment of the present invention.

DETAILED DESCRIPTION OF THE

5

PRESENTLY PREFERRED EXEMPLARY EMBODIMENT

For brevity and ease of explanation, only those aspects of the XML language relevant to explaining how to make and use the present invention are described herein.

Additional details on XML may be found, e.g., in the XML Specification 1.0 found at, *inter alia*, <http://www.w3.org/TR/1998/REC-xml-19980210> and incorporated herein by reference.

10

Of course, the present invention is not limited to the use of XML, and any other programming language having suitable characteristics as described below, e.g., SGML, may be used in its place.

Further, those skilled in the art will readily understand that the present invention is preferably implemented as software executed by multiple processors in a
15 networked computer system which causes the computers to generate appropriate electrical signals as is known in the art. For ease of understanding, the preferred embodiment will be explained with a focus on the software processes and the data they manipulate, rather than on the networked hardware and electrical signals themselves.

20

In a preferred embodiment of the present invention, the analog of the Linda tuple is called an entry, and entries populate entry spaces (similar to tuple spaces). An entry is implemented as an instance of an XML DTD and has the general format

25

```
<entry-name>  
  <field-name-1> field-value-1 </field-name-1>  
  <field-name-2> field-value-2 </field-name-2>  
  .  
  .  
</entry-name>
```

For example, an entry in an embodiment of the present invention used, e.g., in a clothing sales system might look like

```

5      <lot-available>
        <item> sweater </item>
        <fabric> wool </fabric>
        <size> xxl-tall </size>
      </lot-available>

```

10

A field in an entry might have a nested structure:

```

      <lot-available>
        <item> sweater </item>
15      <fabric> wool </fabric>
        <size> <waist> 34 </waist>
          <inseam> 33 </waist> </size>
      </lot-available>

```

15

20

The operations which may be performed on an XML Space are:

```

      write(entry)
      read(template)
      read-if-exists(template)
      take(template)
25     take-if-exists(template)

```

25

Here,

-- write(entry) writes an entry to the entry-space;

-- read(template) performs a blocking, non-destructive copy of an entry

30 matching template (templates are discussed in greater detail below); read-if-

exists(template) performs a non-blocking, non-destructive copy of an entry matching template;

-- take(template) performs a blocking, destructive copy of an entry matching template; and take-if-exists(template) performs a non-blocking, destructive copy of an
 35 entry matching template.

A template is a generalization of an entry in which some of the tagged fields may be under-specified or unspecified. For example, a software entity wishing to locate

available lots of woolen sweaters in any size in the above-described embodiment might call
read-if-exists(template) with the following template:

```
5      <lot-available>
        <item> sweater </item>
        <fabric> wool </fabric>
        <size> * </size>
      </lot-available>
```

- 10 The wildcard * in a field will match any specific value in an analogous field in the XML
Space. More precise means of pattern-matching such as regular expressions can also be used,
as will be readily apparent to those skilled in the art.

The example above was given in connection with articles of clothing for
explanation purposes; however, it will be apparent to those skilled in the art that the
15 applications of entries according to the present invention are numerous if not limitless. For
example, entries can represent any sort of networked or network-proxied resource, object, or
service. An embodiment could provide agent synchronization services, electronic resource-
transfer services, notification services, or virtually anything having to do with distributed
coordinating entities that could make use of XML data.

- 20 Once entries are in an entry space, there must be some means for easily,
quickly and efficiently accessing them. Such techniques according to the preferred
embodiment of the present invention are based upon entry types. The type of an entry is a
means of describing the entry with some degree of generalization and can be implemented in
various ways. For example, the type of an entry can be considered to be characterized by its
25 DTD, or more restrictively by the particular nested structure of its XML tags (without regard
to field values). The latter characterization is more restrictive in the sense that two entries
which have the same type with respect to the latter characterization will also have the same
type with respect to the former characterization, but not necessarily vice-versa. The preferred
choice of representation will be application-dependent and apparent to those skilled in the art.

One technique for addressing entries in an entry space also provides a straightforward means of aggregating disparate entry spaces. It advantageously recognizes that the number of different types of entries in an aggregated entry space will scale sub-linearly with the total number of entries, as shown conceptually in FIG. 1. This basically means that the number of different types of things about which networked communities of entities will want to communicate grows more slowly, in the computational-complexity sense, than the total volume of things. Therefore, doubling (for example) the number of things in an entry space may only increase the number of distinct types by a logarithmic factor or even less. The exact rate of increase will, of course, be application-dependent.

For example, in the clothing example developed above, there are basically a limited number of types of clothing, e.g., shirts, pants, socks, etc. If the number of entries in an entry space was doubled, the number of entry types is likely to increase little, if at all, since most if not all of the new entries will represent a type of clothing probably already present in the entry space.

This principle can be leveraged by implementing a networked hierarchy (called an entry universe) of increasingly condensed representations of sets of entry types present in an entry space as shown in FIG. 2. If an entity, e.g., a process, is seeking a particular entry type, and an entry of that type is not present in a given entry space, the entry universe is traversed upwards (in the direction of increased generality) to a higher-order node called a metaspace containing a generalized description of the nodes below it to seek a group of entry spaces in which an entry of the desired type resides. Then, from that point the entry universe is traversed downwards (in the direction of increased specificity) to find the particular entry space and entry of interest.

As seen in the Figure, leaf nodes in the entry universe are entry spaces containing entries of interest to various networked entities. Each leaf node is connected to

a metaspace 50 in the next higher layer of the entry universe 10. Each metaspace 50 itself is (for first-level metaspaces 50) a tuple-space comprising a collection of pairs each mapping a particular child entry space 30 to a condensed representation of the set of entry types 20 that it contains, or (for higher-level metaspaces) a tuple-space comprising a collection of pairs each mapping a particular metaspace 50 in the next lower level to a condensed representation of the set of entry types 20 which it references. Since the function of the metaspaces is to support aggregation of the XML spaces and not to support explicit XML space retrievals, the metaspaces need not be implemented in XML, and various appropriate programming paradigms will be readily apparent to those skilled in the art.

10 The condensed representations of sets of entry types 20 in the entry spaces 30 are created in a fashion analogous to the creation of digital signatures used for document retrieval. In this latter process, a set of terms (words or other important syntactic elements) is extracted from a document and each extracted term is hashed into a fixed-length bit-vector known as a signature. The signatures for all the terms are superposed and bitwise-OR'ed together to form a signature for the document. Document signatures are stored for each member of a corpus of documents, thus representing essentially a (lossy) compressed version of that corpus. For further information, see, e.g., Sun et al., "Searching the World-Wide Web Using Signature Files", incorporated herein by reference.

20 In order to retrieve a document containing a particular term, the digital signature of that term is created, and compared with the stored document signatures. Any stored signature which "covers" the bits in the term-signature (i.e., has bits set for all the positions where the term-signature does) is a high-probability candidate for containing the term. The term is then compared against the document itself to verify its presence.

25 In the preferred embodiment of the present invention, the "document" role is played by the collection of entries 40 contained within a given entry space 30 (in first-level

metaspaces 50) or by the collection of condensed entry types 20 in metaspaces 50 on the next lower level, and the "terms" which are hashed into signatures are the entries 40 or entry types 20. Thus, to find an entry space 30 containing a particular entry type 20, that entry type 20 is first hashed into a signature which is then compared with the signatures stored in the
5 metaspaces 50. A metaspaces 50 with a signature subsuming the bits of the entry type 20 is a high probability candidate for containing that type 20.

Metaspaces 50 aggregate to form higher-level metaspaces 50 in the entry universe 10 in the following manner. Nominally, one might create the composite signature for a collection of metaspaces 50 by simply superposing the signatures of each member of the
10 collection. From a practical standpoint, this would produce an increasing probability of false hits when searching (via signatures) for entry spaces or metaspaces 50 which contain a particular entry type 20.

To maintain a constant false-hit rate higher in the entry universe 10, the length of a signature in a metaspaces 50 is positively correlated to its level. That is, the signature
15 length of a metaspaces 50 closer to the top entry universe 10 node is longer than that of a metaspaces 50 farther away from the top entry universe 10 node. Thus, all the signatures in a given metaspaces 50 should be converted into longer signatures in an information-preserving fashion. Essentially, the signatures need to be re-sampled, and the most sound means to do this is by mapping them via Fourier methods into the spatial-frequency domain. There, the
20 transforms of the signatures are compressed into a smaller range of spatial frequencies, and then inverse Fourier transforms are computed based on target bit-vectors which are longer than the original ones. Once this has been accomplished, the signatures for all of the metaspaces 50 are superposed to produce a signature for use in the next higher metaspaces 50. This process repeats recursively upwards to the top of the hierarchy. In this way, it is seen that a rate of
25 sampling the mapped descriptions is dynamically determined based on characteristics of the

entry universe. Also, although the Fourier technique is widely known in the art, other appropriate techniques such as simple resampling can also be used.

The particular expansion factor for the signature width at each level of the entry universe 10 depends on the growth in the number of different types of entries 30 represented within a metaspace 50. The optimal expansion factor could be determined dynamically based on the characteristics of the collection at any particular time.

It is not necessary for all immediate children of any particular metaspace 50 in the entry universe 10 to pass upward signatures of the same size. The metaspace 50 may normalize all of the received signatures to be of the same length, for composition of a signature to pass up to its parent, via the spatial frequency methods described above. It is only necessary, when handling searches for particular entry types propagating up the entry universe 10, that each individual entry signature which is propagated be normalized in the same way as was the signature of the corresponding child metaspace 50.

Using this technique for accessing entry spaces 30, one can see how it easily lends itself to the aggregation of disparate spaces. For example, each of the entry spaces 30 may be on a separate networked computer system within the entry universe 10, or several may be resident on different systems. Due to the hierarchical nature of the entry universe 10 and the transformation of element type signatures between levels, the aggregated spaces appear as a single homogenous entry space to a process or other entity accessing it. This further increases the flexibility and wide applicability of the entry space paradigm.

The present invention has been described above in connection with a preferred embodiment thereof; however, this has been done for purposes of illustration only, and the invention is not so limited. Indeed, variations of the invention will be readily apparent to those skilled in the art and also fall within the scope of the invention.

WHAT IS CLAIMED IS:

1. A method of sharing data between a plurality of processors in communication with one another comprising:

5 using a first computer to generate an entry using an extensible markup language, the entry comprising a plurality of fields, each of the fields having a value;
using a computer to store the entry in an entry universe; and
using a second computer to read the entry from the entry universe.

2. The method of claim 1, wherein the second computer reads the entry from
10 the entry universe using a template matching the entry.

3. The method of claim 1, wherein portions of the entry universe are resident on a plurality of computers in communication with one another.

15 4. The method of claim 1, wherein the entry universe comprises:
a plurality of entry spaces; and
a metaspace, the metaspace containing a generalized description of entries in the entry spaces.

20 5. The method of claim 4, further comprising:
using a computer to generate an entry type for entries in one of the plurality of entry spaces;
using a computer to generate a signature for entry types of the entry space; and
using a computer to associate the signature with the metaspace.

25

6. The method of claim 1, wherein the entry universe comprises:

a plurality of entry spaces; and

a plurality of metaspaces;

wherein each one of a first group of the metaspaces is associated with multiple

5 ones of the plurality of entry spaces and contains a generalized description of entries in its associated entry spaces; and

each one of a second group of the metaspaces is associated with multiple ones of the plurality of first group of metaspaces and contains a generalized description of descriptions in its associated first group metaspaces.

10

7. The method of claim 6, wherein the generalized description in each second group metaspace is longer than the generalized description in each of its associated first group metaspaces.

15

8. The method of claim 7, further comprising:

using a computer to Fourier map generalized descriptions of first group metaspaces associated with a second group metaspace to a spatial frequency domain;

using a computer to sample the mapped descriptions into a smaller range of spatial frequencies than in the mapped description;

20

using a computer to inverse map the sampled descriptions; and

using a computer to superpose the inverse mapped descriptions to obtain the generalized description of the second group metaspace.

25

9. The method of claim 8, wherein a rate of sampling the mapped descriptions is dynamically determined based on characteristics of the entry universe.

10. The method of claim 8, further comprising using a computer to normalize generalized descriptions of the first group of metaspaces associated with the second group metaspaces.

5

11. A method of accessing an entry in an entry universe comprising:
causing a computer to access an entry space in the entry universe using a template;

when the entry space does not contain an entry matching the template, causing
10 a computer to access a first metaspaces associated with the entry space and containing generalized descriptions of entry spaces associated therewith;

when the template matches a generalized description in the first metaspaces,
accessing an entry space corresponding to the generalized description; and

when the template does not match a generalized description in the first
15 metaspaces, accessing a second metaspaces containing a generalized description of the first metaspaces.

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 01/2204

A. CLASSIFICATION OF SUBJECT MATTER		
IPC(7) : G06F 15/16		
US CL : 709/201, 202		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols) U.S. : 709/201, 202		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) Please See Continuation Sheet		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	Rowston et al. "BONITA: A set of tuple space primitives for distributed coordination", IEEE System Sciences Proceeding 1997.	1-11
Y, P	US 5,974,420 A (LEHMAN et al.) 26 Oct. 1999 (26.10.1999), abstract	1-11
Y, P	Widergren et al. "XML for Data Exchange", Power Engineering Society Meeting, July 1999, IEEE pp.804-842.	1-11
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.		
* Special categories of cited documents:		
"A"	document defining the general state of the art which is not considered to be of particular relevance	"1"
"E"	earlier application or patent published on or after the international filing date	"X"
"L"	document which may throw doubts on priority claims) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y"
"O"	document referring to an oral disclosure, use, exhibition or other means	
"P"	document published prior to the international filing date but later than the priority date claimed	"E"
Date of the actual completion of the international search 03 August 2000 (03.08.2000)		Date of mailing of the international search report 22 AUG 2000
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703)305-3230		Authorized officer D. Dini <i>James R. Matthews</i> Telephone No. 305 9600

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US96/12201

Continuation of B. FIELDS SEARCHED Item 3: IEEE
USPAT

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.